

PDSF Users Meeting

- * PDSF performance
- * announcements
- * AOB

April 11, 2017

Jan Balewski

aggregated load on PDSF interactive nodes

<https://portal-auth.nersc.gov/pdsf-mon/>

Table 2 List of big users

user	use (core)	num jobs	use CPU (day)	list of jobs	flag
whl6	0.9	1	3.51	['pdsf7:emacs']	1
rlinehan	2.0	2	0.87	['pdsf8:BACCARATExecuta', 'pdsf8:BACCARATExecuta']	0
pchu	2.0	2	26.90	['pdsf7:emacs', 'pdsf8:emacs']	1
tdavison	1.0	1	5.24	['pdsf7:BACCARATExecuta']	1


some programs
tend to stuck - kill them
before log-out

Query MongoDB : **PLOT history** last 502 records, set to last: [day](#), [week](#), [month](#)

Change time zone for x-axis : **UTC off** , **UTC on**


PLOT:

used cores, argegated over pdsf6,7,8




CPU Utilization is not great

- not enough jobs submitted and/or
- few STAR user submit enough jobs but they are throttled because xrootd can't keep up w/ job demand



UGE queue load

Not enough submitted jobs
or by users w/ reduced running privileges



Utilization of PDSF compute nodes

April average

Served 3364 vCores
Capacity 4000 vCores
CPU utilization 0.76

end-time was set to 2017-04-10 20:15, DB collName=u

user \ timeRange	sum last 7_days
0-all-users : owner	59.2 (cpu*year) , 235462 jobs wallT fract=1.000, CPU eff=0.76
alice : project	35.9 (cpu*year) , 43544 jobs wallT fract=0.607, CPU eff=0.63
atlas : project	2.1 (cpu*year) , 34038 jobs wallT fract=0.036, CPU eff=0.98
dayabay : project	0.6 (cpu*year) , 1960 jobs wallT fract=0.009, CPU eff=0.92
dybprod : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00
hack : owner	0.3 (cpu*year) , 1447 jobs wallT fract=0.006, CPU eff=0.95
lux : project	0.0 (cpu*year) , 21 jobs wallT fract=0.000, CPU eff=0.94
lz : project	0.0 (cpu*year) , 3600 jobs wallT fract=0.000, CPU eff=0.49
majorana : project	0.2 (cpu*year) , 9178 jobs wallT fract=0.003, CPU eff=0.77
star : project	20.2 (cpu*year) , 140159 jobs wallT fract=0.341, CPU eff=0.92
staremb : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00
zmarshal : owner	0.0 (cpu*year) , 7 jobs wallT fract=0.000, CPU eff=0.15

sum last 30_days
276.5 (cpu*year) , 1140416 jobs wallT fract=1.000, CPU eff=0.76
129.2 (cpu*year) , 167883 jobs wallT fract=0.467, CPU eff=0.68
7.8 (cpu*year) , 109082 jobs wallT fract=0.028, CPU eff=0.76
18.9 (cpu*year) , 34902 jobs wallT fract=0.068, CPU eff=0.35
0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00
18.7 (cpu*year) , 34377 jobs wallT fract=0.068, CPU eff=0.34
1.5 (cpu*year) , 18492 jobs wallT fract=0.005, CPU eff=0.62
0.1 (cpu*year) , 5581 jobs wallT fract=0.000, CPU eff=0.82
12.0 (cpu*year) , 238318 jobs wallT fract=0.043, CPU eff=0.74
106.3 (cpu*year) , 562115 jobs wallT fract=0.385, CPU eff=0.92
14.5 (cpu*year) , 8624 jobs wallT fract=0.052, CPU eff=1.00
0.0 (cpu*year) , 31 jobs wallT fract=0.000, CPU eff=0.19

March Average

Served 3460 vCores
Capacity 4000 vCores
CPU utilization 0.71

sum last 30_days
282.4 (cpu*year) , 1760555 jobs wallT fract=1.000, CPU eff=0.71
114.1 (cpu*year) , 202731 jobs wallT fract=0.404, CPU eff=0.82
3.7 (cpu*year) , 110413 jobs wallT fract=0.013, CPU eff=0.68
63.8 (cpu*year) , 101074 jobs wallT fract=0.226, CPU eff=0.17

1.8 (cpu*year) , 8886 jobs wallT fract=0.007, CPU eff=0.17
5.8 (cpu*year) , 381699 jobs wallT fract=0.020, CPU eff=0.95
9.5 (cpu*year) , 173790 jobs wallT fract=0.034, CPU eff=0.93
82.1 (cpu*year) , 780320 jobs wallT fract=0.291, CPU eff=0.90

I/O monitoring added


added dybspade

UGE usage report cron-generated: 2017-04-12_13.33, took 26.6 minutes
end-time was set to 2017-04-12 13:07, DB collName=ugeJobs2017

user \ timeRange	sum last 7_days	sum last 3_days	sum last 30_days
0-all-users : owner	56.7 (cpu*year) , 254209 jobs wallT fract=1.000, CPU eff=0.75	18.0 (cpu*year) , 112571 jobs wallT fract=1.000, CPU eff=0.68	271.2 (cpu*year) , 1157182 jobs wallT fract=1.000, CPU eff=0.75
alice : project	31.7 (cpu*year) , 61204 jobs wallT fract=0.558, CPU eff=0.57	7.3 (cpu*year) , 34422 jobs wallT fract=0.405, CPU eff=0.26	125.9 (cpu*year) , 188688 jobs wallT fract=0.464, CPU eff=0.65
atlas : project	1.4 (cpu*year) , 28116 jobs wallT fract=0.025, CPU eff=0.98	0.0 (cpu*year) , 4086 jobs wallT fract=0.002, CPU eff=0.95	6.5 (cpu*year) , 107109 jobs wallT fract=0.024, CPU eff=0.81
dayabay : project	0.6 (cpu*year) , 1995 jobs wallT fract=0.011, CPU eff=0.94	0.1 (cpu*year) , 373 jobs wallT fract=0.006, CPU eff=0.95	17.6 (cpu*year) , 33023 jobs wallT fract=0.065, CPU eff=0.36
dybprod : owner	0.0 (cpu*year) , 10 jobs wallT fract=0.000, CPU eff=0.04	0.0 (cpu*year) , 10 jobs wallT fract=0.000, CPU eff=0.04	0.0 (cpu*year) , 10 jobs wallT fract=0.000, CPU eff=0.04
dybspade : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00
hack : owner	0.4 (cpu*year) , 1653 jobs wallT fract=0.007, CPU eff=0.95	0.1 (cpu*year) , 363 jobs wallT fract=0.006, CPU eff=0.95	17.4 (cpu*year) , 32492 jobs wallT fract=0.064, CPU eff=0.35
lux : project	0.0 (cpu*year) , 141 jobs wallT fract=0.000, CPU eff=1.00	0.0 (cpu*year) , 120 jobs wallT fract=0.002, CPU eff=1.00	1.4 (cpu*year) , 18211 jobs wallT fract=0.005, CPU eff=0.60
lz : project	0.0 (cpu*year) , 3561 jobs wallT fract=0.000, CPU eff=0.46	0.0 (cpu*year) , 2511 jobs wallT fract=0.001, CPU eff=0.37	0.1 (cpu*year) , 5581 jobs wallT fract=0.000, CPU eff=0.82
majorana : project	0.2 (cpu*year) , 10404 jobs wallT fract=0.004, CPU eff=0.75	0.1 (cpu*year) , 5051 jobs wallT fract=0.003, CPU eff=0.70	11.7 (cpu*year) , 216786 jobs wallT fract=0.043, CPU eff=0.74
star : project	22.6 (cpu*year) , 145961 jobs wallT fract=0.398, CPU eff=0.93	10.4 (cpu*year) , 65848 jobs wallT fract=0.576, CPU eff=0.91	107.3 (cpu*year) , 583663 jobs wallT fract=0.396, CPU eff=0.92
staremb : owner	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00	0.0 (cpu*year) , 0 jobs wallT fract=0.000, CPU eff=0.00	14.5 (cpu*year) , 8624 jobs wallT fract=0.053, CPU eff=1.00
zmarshal : owner	0.0 (cpu*year) , 7 jobs wallT fract=0.000, CPU eff=0.14	0.0 (cpu*year) , 3 jobs wallT fract=0.000, CPU eff=0.11	0.0 (cpu*year) , 31 jobs wallT fract=0.000, CPU eff=0.20

UGE CPU utilization


116 UGE nodes, 4578 job slots (potentially),
1300 jobs run, cluster is 40% loaded (last Monday)



More details: <http://portal.nersc.gov/project/mpccc/balewski/ugeLoad/latest/>

RAM/job distribution

RAM/job < 0.5 GB



xrootd files re-distribution





Table produced on 2017-04-10 18:53

highDiscFract		lowFreeDisk	
38 nodes		38 nodes	27 nodes

/project utilization

<http://portal.nersc.gov/project/star/jthaeder/diskUsage/overview/indexExt.html>

STAR is out of
comfort zone

FillStatus (Quota): *PROJECT* (2017-04-10 20:04)

star - size

64.526/70.000 TB (92.18%)

star - inodes

19891240/200000000 (99.45%)

starprod - size

118.555/130.000 TB (91.19%)

starprod - inodes

8926036/200000000
(44.63%)

alice - size

42.329/60.976 TB (69.41%)

alice - inodes

16694712/250000000 (66.77%)

FillStatus (Quota): *PROJECTA* (2017-04-10 20:04)

starprod - size

168.418/190.000 TB (88.64%)

starprod - inodes

6112489/200000000
(30.56%)

Announcements

Bi-weekly office hours 12:30 -2:30pm

Thursday, April 13, 27, 59-4016-CR

PDSF user meeting

- Tuesday, May 9, 11am - 12pm, 59-3034-CR


Outages

- Outage to add new Cori Haswell Cabinets: second half of April
- Outage for Quarterly Maintenance: May 16–18
- Memorial Day Holiday (No Consulting or Account Support): May 29

Devaluation:

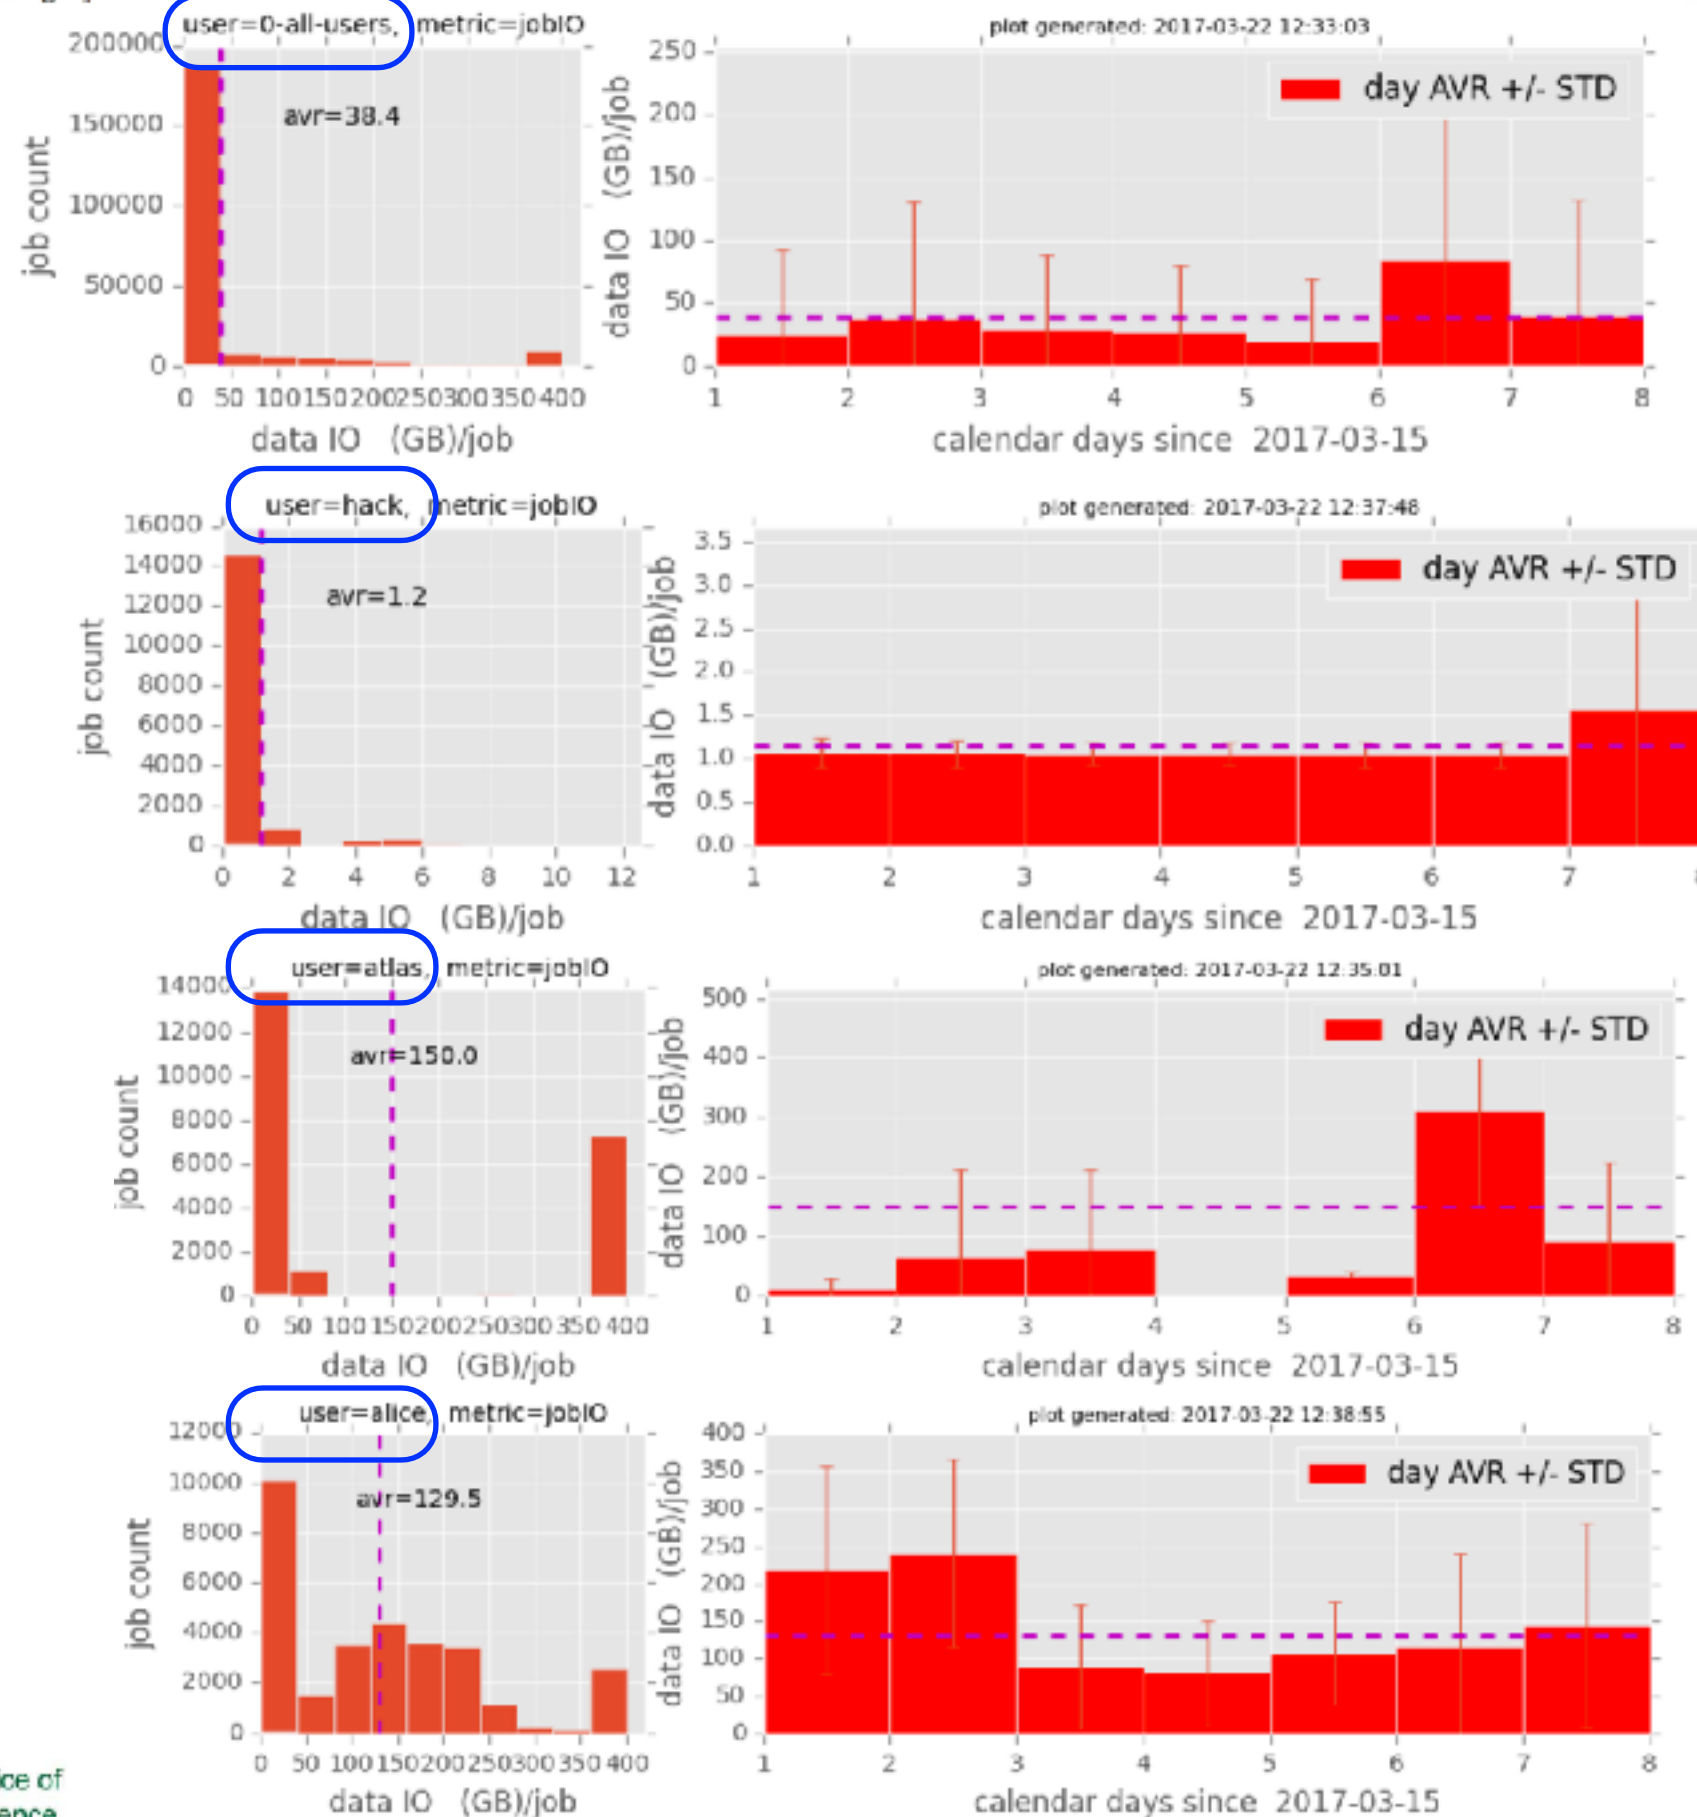
First Quarterly Allocation Reduction: April 12 (this week)

job io monitor



job IO - week average

Fig.3 jobIO



STAR embedding on Cori in Shifter/SL64

60 r4s task per node

free -g

	total	used	free	shared	buffers	cached
Mem:	126	73	52	0	0	10
-/+ buffers/cache:		62	63			
Swap:	0	0	0			

top - 23:20:07 up 3 days, 9:24, 0 users, load average: 59.52, 59.48, 54.48

Tasks: 677 total, 62 running, 615 sleeping, 0 stopped, 0 zombie

%Cpu(s): 41.1 us, 0.2 sy, 0.0 ni, 58.6 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st

KiB Mem: 13216380+total, 76871760 used, 55292052 free, 100120 buffers

KiB Swap: 0 total, 0 used, 0 free. 10958920 cached Mem



PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND		
1	16995	balewski	20	0	1484276	1.074g		93276	R	100.00	0.852	39:53.75	root4star
2	17777	balewski	20	0	1495544	1.095g		97052	R	100.00	0.869	38:44.61	root4star
3	20928	balewski	20	0	1458492	1.051g		95424	R	100.00	0.834	39:06.72	root4star
4	24046	balewski	20	0	1477820	1.070g		96360	R	100.00	0.849	39:15.17	root4star
5	26774	balewski	20	0	1544608	1.142g		93240	R	100.00	0.906	38:21.16	root4star
6	9801	balewski	20	0	1477340	1.070g		95764	R	100.00	0.849	40:12.45	root4star
7	9802	balewski	20	0	1513796	1.109g		98040	R	100.00	0.880	39:59.50	root4star
8	10072	balewski	20	0	1514864	1.107g		98828	R	100.00	0.878	40:15.63	root4star
9	10353	balewski	20	0	1500460	1.093g		97844	R	100.00	0.867	40:33.31	root4star
10	10732	balewski	20	0	1474472	1.069g		96984	R	100.00	0.848	40:12.61	root4star
11	11126	balewski	20	0	1464792	1.055g		94736	R	100.00	0.837	34:58.26	root4star
12	11516	balewski	20	0	1467348	1.055g		93572	R	100.00	0.837	32:38.22	root4star
13	11907	balewski	20	0	1489552	1.082g		98036	R	100.00	0.859	39:49.53	root4star
14	12297	balewski	20	0	1463256	1.056g		96176	R	100.00	0.838	37:36.07	root4star
15	12688	balewski	20	0	1500488	1.098g		97752	R	100.00	0.871	40:12.60	root4star
51	28717	balewski	20	0	1475208	1.066g		95408	R	100.00	0.846	37:34.33	root4star
52	29106	balewski	20	0	1529152	1.122g		94064	R	100.00	0.890	38:09.74	root4star
53	29494	balewski	20	0	1490808	1.086g		96168	R	100.00	0.861	36:11.60	root4star
54	29889	balewski	20	0	1500088	1.097g		97308	R	100.00	0.870	37:53.78	root4star
55	30337	balewski	20	0	1504052	1.095g		95376	R	100.00	0.869	37:53.20	root4star
56	30669	balewski	20	0	1480480	1.075g		97076	R	100.00	0.853	38:39.54	root4star
57	31058	balewski	20	0	1449216	1.038g		93300	R	100.00	0.824	37:20.52	root4star
58	31446	balewski	20	0	1484268	1.073g		92828	R	100.00	0.852	38:02.71	root4star
59	31926	balewski	20	0	1461488	1.054g		96168	R	100.00	0.836	37:08.48	root4star
60	32224	balewski	20	0	1463356	1.056g		96260	R	100.00	0.838	37:36.12	root4star

STAR embedding on Cori in Shifter/SL64

600 r4s 2-hour jobs starting within few seconds

```
bfcMixer_Tpx($NUM_EVE, "$daqN", "$tagsN", 0, 12.0, -0.8, 0.8, -200.0, 200.0, 100.0, 169, 8, triggers, "P16idpp200", "FlatPt", 0, "$fzdN");
```

Some r4s tasks spent 30 minutes in 1st event
this is 30% tax on 2hour task



Event [no. 1/run 16082055/evt. 1514/Date.Time 20150324.25959/sta 0]	Real Time =	2007.52 seconds	Cpu Time =	140.38 seconds
Event [no. 2/run 16082055/evt. 8718/Date.Time 20150324.30002/sta 0]	Real Time =	95.34 seconds	Cpu Time =	94.98 seconds
Event [no. 3/run 16082055/evt. 29632/Date.Time 20150324.30010/sta 0]	Real Time =	87.19 seconds	Cpu Time =	86.98 seconds
Event [no. 4/run 16082055/evt. 29957/Date.Time 20150324.30010/sta 13]	Real Time =	30.59 seconds	Cpu Time =	30.50 seconds
Event [no. 5/run 16082055/evt. 41749/Date.Time 20150324.30015/sta 13]	Real Time =	35.47 seconds	Cpu Time =	35.40 seconds
Event [no. 6/run 16082055/evt. 46231/Date.Time 20150324.30017/sta 0]	Real Time =	111.63 seconds	Cpu Time =	111.14 seconds
Event [no. 7/run 16082055/evt. 53079/Date.Time 20150324.30019/sta 13]	Real Time =	51.60 seconds	Cpu Time =	51.44 seconds